Warning Concerning Copyright Restrictions

The Copyright law of the United States (Title 17, United States Code) governs the making of photocopies or other reproductions of copyright material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction not be "used for any purposes other than private study, scholarship, or research." If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of "fair use," that user may be liable for copyright infringement.

CHAPTER 9

INCIDENTAL FORGETTING

Michael C. Anderson

Over the Christmas holiday, my sister asked, "Do you remember when you knocked over the Christmas tree?" I said, "What are you talking about? I never did that!" Puzzled, my sister said, "Yes you did, don't you remember?" My brother added, "Yes, you were hurrying to squeeze behind the tree so you could take a picture of Aunt Dottie and Uncle Jim as they came up the driveway when you knocked the tree over." Indignant, I said, "What ... what are you talking about ... you must be mixing me up with someone else." My father insisted, "No, you definitely knocked the tree over. It was a big mess, and we made fun of you for it." He added that he remembered me feeling bad about ruining the tree, even though everyone said it was okay. They simply couldn't believe that I had forgotten this.

Reluctantly, I accepted that this event must have happened. I struggled to recall details and couldn't come up with anything. I said, "When did this happen? When I was a kid?" My sister replied, "No, it was about 3 or 4 years ago when we were in New York." I was shocked. I called my other brother and he confirmed every detail and was able to recall the year it had occurred. In fact, I remembered that Christmas in New York and my mother's new camera (which I was using), but I simply could not remember this event. After many months and repeated searching, I still could not bring any trace of the experience to mind.

From life's embarrassing mishaps, to the mundane details of our daily life, many of our memories are forgotten. How and why are certain memories lost while others remain vivid for a lifetime? © Lambert/Archive Photos.

Before you start wondering whether I'm amnesic, consider how much of *your* life *you* remember. Take a break from reading and try an exercise. Get out a sheet of paper and list everything that you did from the time you got up until the time you went to bed yesterday, including details about who you saw, and any conversations or thoughts. Chances are, you did pretty well and came up with a lot of detail. Perhaps you left out one or two minor things that you would recall if reminded. Next, do the same thing for the day that occurred *1 week* earlier. You can probably still recall a lot, but with much more effort, and you most likely feel that you are forgetting more. Finally, try the same thing, but for a day that occurred exactly *1 year* prior to yesterday. Try very hard. Most likely, after significant effort, you probably didn't recall much except perhaps some broad outlines that you are probably only guessing at, and only then after much reconstruction. The same uncomfortable fact is true for the majority of the days in your life, except for truly special events and the recent past.

In fact, consider this: *this very moment* that you are consciously experiencing, will, if your history serves as any guide, join the rest of those lost experiences. One cannot help but wonder how it is possible for something that is the full focus of your consciousness right now to ultimately be so completely lost. Is this the fate of all experience? When you are 80, will you only remember 1% of your life in any detail? Are all of your memories there, and just inaccessible?

The function of memory is never more conspicuous and astonishing than when it fails us. In this chapter, we consider the mechanisms that underlie forgetting. One might wonder why forgetting should be treated in a separate chapter from retrieval, in which we discussed why retrieval fails us. Indeed, retrieval failure *is* a form of forgetting. Forgetting is worthy of being distinguished, however, because of the potential for distinct forgetting processes that contribute to retrieval failure. Moreover, an emphasis on forgetting leads one to focus on changes in retrievability over time. What factors produce those changes? What would life be like if we never forgot?

In addressing these questions, research on memory has focused on both **incidental forgetting** and **motivated forgetting**. Incidental forgetting occurs without the intention to forget; motivated forgetting occurs when people engage processes or behaviors that intentionally diminish accessibility for some purpose. It is likely that to explain the full range of experiences that people have with forgetting, theories of both types of forgetting are needed. We discuss incidental forgetting here, and motivated forgetting in the next chapter.

A REMARKABLE MEMORY

What would it be like to remember every thing that ever occurred to you? Although no such person has yet been found, there are people with astounding memory. Recently, Elizabeth Parker, Larry Cahill, and James McGaugh (2006) reported the fascinating case of AJ, a 41-year-old woman who had a breathtaking capacity to remember her past. AJ remembers every single day of her life since her teens in extraordinary detail. Mention any date over several decades and she finds herself back on that day, reliving events and feelings as though they happened yesterday. She can tell you what day of the week it was, events that took place on all surrounding days, and intricate details about her thoughts, feelings, and public events, all of which can be verified by personal diaries she has kept over 30 years. AJ reports that these memories are vivid, like a running movie, and full of emotion. Her remembering feels automatic, and not under conscious control, a claim supported by the fact that her recollections occur immediately, with no struggle.

One might think that having such a remarkable memory would be wonderful. But it's not all good. When unpleasant things happen, AJ wishes she could forget, and the constant bombardment by reminders is distracting and sometimes troubling. In AJ's words:

> *My memory has ruled my life ... It is like my sixth sense ... There is no effort to it ...*

KEY TERMS

Incidental forgetting: Memory failures occurring without the intention to forget.

Motivated forgetting: A broad term encompassing intentional forgetting as well as forgetting triggered by motivations, but lacking conscious intention.

> *I want to know why I remember everything. I think about the past all the time ... It's like a running movie that never stops. It's like a split screen. I'll be talking to someone and seeing something else ... Like we're sitting here talking and I'm talking to you and in my head I'm thinking about something that happened to me in December 1982, December 17th, 1982, it was a Friday, I started to work at Gs [a store] ... I only have to experience something one time and I can be totally scarred by it ... I can't let go of things because of my memory ... Happy memories hold my head together ... I treasure these memories, good and bad ... I can't let go of things because of my memory, it's part of me ... When I think of these things, it is kind of soothing ... I knew a long time ago, I had an exceptional memory ... I don't think I would never want to have this but it's a burden.*

Parker et al. (2006) have termed AJ's condition *hyperthymestic syndrome*, from the Greek word *thymesis*, meaning remembering. In short, AJ has uncontrollable remembering. Clearly, AJ's experience of life is very different from ours and illustrates a cost she pays for her perfect memory: she can remember the good times but at the cost of suffering the persistence of bad times. Would you choose AJ's memory over your own? Perhaps forgetting is not all bad. Later in this chapter, we will discuss the possibility that forgetting serves a useful function.

THE FUNDAMENTAL FACT OF FORGETTING

Clearly, AJ's experience is atypical, as most of us forget the events of our days. How are we to understand forgetting? A good place to begin discussing this phenomenon is to acknowledge a fundamental fact: for most people (and organisms), *forgetting increases as time progresses.* Although this surely comes as no surprise, you might not have considered the nature of the relationship between memory and time. If you had to guess, would you say that people forget at a constant rate? To address this question, one simply needs to measure how likely forgetting is as a memory grows older. Once again, the classic study was conducted by Hermann Ebbinghaus (1913), using himself as the participant and nonsense syllables as the material to be learned. Ebbinghaus learned 169 separate lists of 13 nonsense syllables then relearned each list after an interval ranging from 21 minutes to 31 days. He always found that some forgetting had occurred and used the amount of time required to learn the list again as a measure of how much had been forgotten. He found a clear relationship between time and retention.

You will recall from Chapter 3 that the relationship between learning and remembering was more or less linear, with the long-term memory store behaving rather like a bath being filled by a tap running at a constant rate. But how about forgetting? Is it simply like pulling the plug out of the bath, causing information to be lost at a constant rate, or is the relationship less straightforward? The results obtained by Ebbinghaus are shown in Figure 9.1. This graph represents a quantitative relationship between memory and time, referred to as a **forgetting curve**, or sometimes a **retention function**. As you can see, Ebbinghaus's forgetting was extremely rapid at first, but it gradually slowed down over time; the rate of forgetting he exhibited was more logarithmic than linear. As with Ebbinghaus's other work, this result

KEY TERM

Forgetting curve/retention function: The logarithmic decline in memory retention as a function of time elapsed, first described by Ebbinghaus.

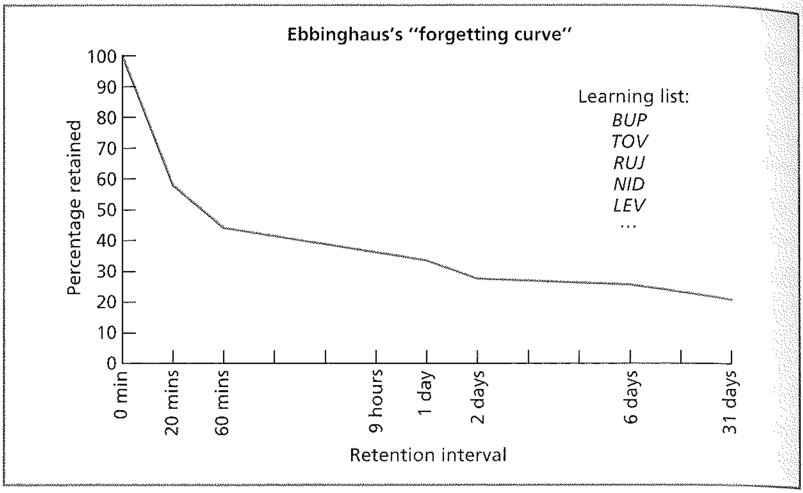

Figure 9.1 The forgetting curve that Ebbinghaus (1913) obtained when he plotted the results of one of his forgetting experiments. His finding, that information loss is very rapid at first and then levels off, holds true for many types of learned materials. Data from Ebbinghaus (1913).

has stood the test of time and applies across a wide range of learning conditions.

Most studies on the rate of forgetting have, like Ebbinghaus's, concerned themselves with highly constrained materials such as lists of nonsense syllables or unrelated words. Is this representative of what happens to personal memories? What happens when more realistic material is recalled over longer intervals? Answering this presents a major problem. Consider the question posed earlier about what you were doing one year ago. If you were to give an answer, how would I know whether you were correct? It is extremely unlikely that the necessary information remains available. One solution is to question respondents about events that were sufficiently noteworthy to attract the attention of most people at the time they happened. This strategy was followed by Meeter, Murre, and Janssen (2005), who selected headlines in both newspapers and television broadcasts for each day over a 4-year period. They amassed over 1000 questions about distinct and dateable events, of which each participant would answer a randomly chosen 40. Cleverly, these investigators used the internet to attract participants, allowing them to test the memory for over 14,000 participants from widely different age groups from countries across the world. They tested their respondents' memory for these events by both recall and recognition.

The results obtained by Meeter et al. (2005) show that substantial forgetting of public events does occur, with participants' recall for the events dropping from 60 to 30% in just a single year. The forgetting curves showed a steep initial decline, followed by a slowed rate of forgetting at longer delays, especially when recall was tested, much like that observed with nonsense syllables by Ebbinghaus over a century ago. They also found that people performed much more poorly when their recall was tested, recalling only 31% of the answers correctly over the years, compared to 52% correct when they simply had to recognize the right answer from among options. These findings lend confidence to the basic conclusions about forgetting from laboratory studies.

The forgetting curves we have discussed so far have been concerned mainly with memory for distinct events, which are relatively poorly learned. What of information that has been more thoroughly and deliberately learned? Light was thrown on this by an intriguing study by Bahrick, Bahrick, and Wittlinger (1975), who traced 392 American high-school graduates and tested their memory for the names and portraits of classmates. Their study showed that the ability to both

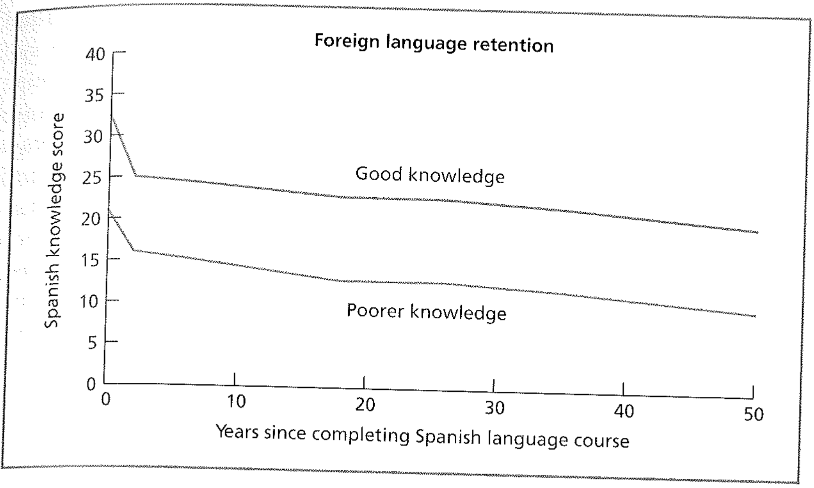

Figure 9.2 People who learned Spanish in college showed rapid forgetting over the first 3 or 4 years, followed by remarkably little forgetting over the next 30 years (Bahrick, 1984). Those who had a good knowledge (upper line, in blue) continued to have a clear advantage, even 50 years later. Data from Bahrick (1984).

recognize a face or a name from among a set of unfamiliar faces or names and to match up names with faces, remained remarkably high for over 30 years. In contrast, the ability to *recall* a name in response to a person's picture showed more extensive forgetting, just as was found in the previously discussed study of memory for major news events.

Harry Bahrick is a Professor at Ohio Wesleyan University, which—in common with many American colleges—has an annual reunion for alumni. Bahrick has made ingenious use of this tradition to study the retention by alumni of a range of material: from the geography of the town where the university is located to the vocabulary of foreign languages learned at college. Figure 9.2 shows the effect of delay on memory for a foreign language (Spanish, in this case). The most striking feature of the graph is the way in which forgetting levels out after about 2 years, with little further loss up to the longest delay, virtually 50 years later. It is as if forgetting occurs only up to a certain point, beyond which memory traces appear frozen. Using an analogy to the permanently frozen ground in Polar regions, known as permafrost, Bahrick (1984) has suggested the term *permastore* for this stable language learning performance. The second point to note is that the overall retention is determined by the level of initial learning, at least as far as learning a foreign language is concerned. Thus, for well-learned materials, it seems, the forgetting curve might flatten out after an initial period of forgetting, and show little additional forgetting over long periods.

ON THE NATURE OF FORGETTING

The discussion of forgetting functions raises an issue concerning what counts as "forgetting." The studies by Meeter et al. and by Bahrick found much greater forgetting when recall was tested, compared to recognition. In fact, this is a robust pattern: *Recognition is generally easier than recall.* A reasonable conclusion to draw from this fact is that recognition tests reveal that more often resides in memory than is measured by recall. If so, is it truly fair to characterize failures to recall information evident in these forgetting functions as actual forgetting, when many of those unrecalled traces reside in memory? Shouldn't we reserve the term *forgetting* to refer to the permanent loss of traces? This issue highlights a distinction aptly named by Endel Tulving—the distinction between a memory's availability in the cognitive system (whether it is in storage or not), and its accessibility (whether one can access a memory, given that it is stored): the **accessibility/availability distinction.** Should we count inaccessibility as forgetting, or only unavailability?

Unfortunately, reserving *forgetting* to refer only to memories made unavailable renders it impossible to ever measure forgetting. The reason is that determining whether a memory has been permanently lost is quite a bit trickier than one might suspect. What will be our evidence of unavailability? Failed recall? Clearly not, as the foregoing results establish. Failed recognition? Again, recognition can fail, even when it can be proved that a trace is in memory, given the proper reinstatement of context. Although an experience might seem lost forever, perhaps the right cue has just not come along. As discussed in Chapter 8, it took the sight of the passport in the box to pry loose my memory for storing it there. It is thus quite difficult to distinguish inaccessibility from unavailability. Moreover, when memories transition from being recallable to only being recognizable, this might, in principle, be due to weakening of the trace. Permanent loss might not be all or none, but might happen in a graded fashion. For these reasons, and because reduced accessibility is a memory failure, inaccessibility is considered forgetting.

FACTORS THAT DISCOURAGE FORGETTING

The studies by Harry Bahrick illustrate how forgetting, although perhaps inevitable for many memories, might be slowed for some types of knowledge. Which factors discourage forgetting? One obvious point is that if you learn something well to begin with, forgetting is less likely, or at least it takes much longer. But are there some ways of strengthening a memory that increase resistance to forgetting more than others? What memories will you have when you are 80?

The apparent flattening-out of the forgetting curve over time demonstrates that memories are not equally vulnerable to forgetting at all points in their history. Another way of describing the relationship of time and memory is in terms of *Jost's Law*, named after a nineteenth-century psychologist, which states that if two memories are equally strong at a given time, then the older of the two will be more durable and forgotten less rapidly. It is as if two opposing forces are at work to determine retention over time; the mechanisms of forgetting, but also some process that makes surviving memories grow tougher with age. Indeed, it is widely believed that new traces are initially vulnerable to disruption until they are gradually stamped into memory. The time-dependent process by which a new trace is gradually woven into the fabric of memory and by which its components and their interconnections are cemented together is known as *consolidation*. At least two types of consolidation have been proposed. According to research on *synaptic consolidation*, the imprint of experience takes time to solidify because it requires structural changes in the synaptic connections between neurons. These modifications rely on biological processes that may take hours to days to complete (Dudai, 2004). Until those structural changes occur, the memory is vulnerable. Research also implicates a process known as *systemic consolidation*, which holds that the hippocampus is initially required for memory storage and retrieval but that its contribution diminishes over time until the cortex is capable of retrieving the memory on its own (Squire, 1992c; Dudai, 2004). As is discussed further in Chapter 11, the hippocampus is thought to accomplish this by recurrently reactivating the

KEY TERM

Accessibility/availability distinction: Accessibility refers to the ease with which a stored memory can be retrieved at a given point in time. Availability refers to the distinction indicating whether a trace is or is not stored in memory.

brain areas involved in the initial experience (e.g. the areas involved in hearing the sounds, seeing the sights, essentially "replaying" the memory) until these areas are interlinked in a way that could recreate the original memory. Until the memory becomes independent of the hippocampus, it is vulnerable to disruption. Estimates of the duration of systemic consolidation vary, with some evidence suggesting that it might take years in humans. So, it seems that a process exists that strengthens memories over time, retarding their forgetting, and that this process involves recurring retrieval of some sort.

Interestingly, intentionally retrieving an experience also has an especially potent effect on the rate at which a memory is forgotten. This fact was illustrated compellingly by Marigold Linton (1975), using herself as a participant. Every day for 5 years, she noted in her diary two events that had occurred. At predetermined intervals she would randomly select events from her diary and judge whether she could recall them. Given the fact that she was sampling in this way any given event could crop up many times. She was therefore able to analyze her results to find out what effect earlier recalls had on the later memorability of the event. Her results are shown in Figure 9.3; the items that were not retested showed dramatic forgetting over a 4-year period (65% forgotten). Even a single test was enough to reduce forgetting, whereas items tested on four other occasions showed an impressively low probability of forgetting after 4 years (only 12% forgotten). So, it seems that personal memories, if retrieved periodically, grow resistant to forgetting, in much the same way as did the cases of permastore for well-learned material reported by Bahrick and colleagues. Other examples of the memory-enhancing power of retrieval are discussed in Chapter 16 on improving your memory.

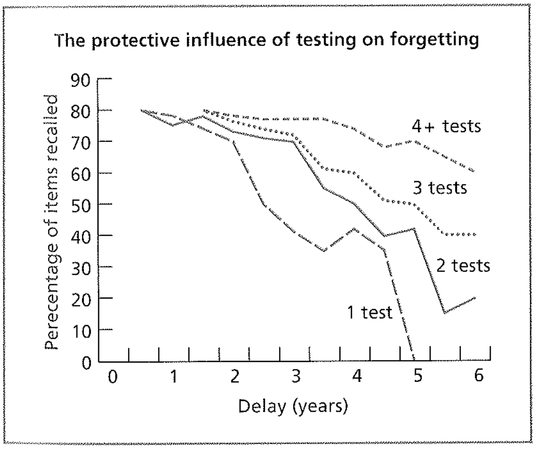

Figure 9.3 The probability of remembering something depends on the number of times it has been retrieved or called to mind. Recalling an event reduces the rate of forgetting. Data from Linton (1975).

Although retrieval enhances retention, we must be cautious about what is being retrieved. People are tempted to assume that if they are recalling something that happened 20 years ago, that they are recalling a 20-year-old memory. This might be true if we have not recalled the memory in the interim. However, if we have retrieved the memory at all, perhaps we are retrieving a memory of what we have retrieved previously. The event of retrieving something is itself a memory, with its own context, and particulars. The more often that we retrieve an experience, the more of these retrieval events will exist in memory. As long as the information retrieved each time is accurate and complete, this process will enhance recall. If recollections are incomplete or inaccurate due to reconstructive inferences, what we remember may not be what originally happened.

It appears, then, that retrieval might play a very special role in determining which elements of experience will be preserved throughout our lives. Each time that we get together and reminisce with friends or family, we are implicitly selecting which memories to more firmly establish. And for those of us who keep a diary, reviewing the day's events and retrieving them not only provides an objective record of their occurrence but might also increase the longevity of those memories, especially if they are reviewed from time to time. Retrieval clearly has a special effect on retention. Later on, I will discuss research demonstrating that, ironically, retrieval also plays a powerful and complementary role in determining what we forget.

FACTORS THAT ENCOURAGE INCIDENTAL FORGETTING

Knowing that retrieval retards forgetting is useful, but why does forgetting occur in the first place? What factors contribute to retention loss? Experimental psychologists have traditionally emphasized incidental forgetting, stressing the involvement of passive processes that occur as a by-product of changes in the world or the person. For example, forgetting has been attributed to decay, contextual shifts, and to **interference**. This passive view fits the general feeling most of us have that we are the unwilling victims of memory loss. This perspective often fits reality: we do forget things unintentionally, even when they are important. Here, we consider several of the most important factors.

Passage of time as a cause of forgetting

The most obvious way of describing the forgetting curve is that memory gets worse as time goes by. Perhaps the cause is that simple: memory traces grow weaker with time. Memories may simply fade, rather like a notice that is exposed to sun and rain gradually fades until it becomes illegible. The idea that memories get weaker over time is known as **trace decay**. Many investigators favor the view that trace decay partially determines the loss of information from verbal and visual working memory (e.g. Broadbent, 1958; Baddeley, 1986; Cowan, 1988; Page & Norris, 1998; Towse, Hitch, & Hutton, 2000; Gold, Murray, Sekuler, Bennett, & Sekuler, 2005), although this approach has its critics (Nairne, 2002). Decay also plays a role in how theorists think about repetition priming and familiarity, with some proposing that these effects decay quickly (e.g. Eichenbaum, 1994; McKone, 1998; Yonelinas & Levy, 2002). Many proposals about trace decay have in common the idea that activation decays gradually, even if the item remains stored. For example, recent exposure to the word *helmet* might activate a pre-existing concept. Although activation might fade, the concept remains.

There is another sense of decay, however, in which a memory's structural elements degrade, not just activation levels. Thus, associations between features or the features themselves might deteriorate. Does this happen? This issue is related to the age-old question of whether memories are stored permanently and merely grow inaccessible. On one level, the answer seems obvious: memories are not permanent and decay must exist. We cannot disregard that we are biological beings. Our memories survive in tissue that continually changes, with neurons dying and connections weakening or being modified. We know, for instance, that a time-dependent process degrades the synaptic connections between neurons that support a recently learned behavior in *Aplysia* (a sea slug), with a corresponding degradation in the learned behavior (Bailey & Chen, 1989). It is not far-fetched to believe that a similar degradation occurs in humans, perhaps underlying time-dependent decay. If neurons die, and connections degrade, the survival of memories over long stretches of time in fact seems the greater mystery.

Although decay seems inevitable, experimental psychologists are rightly skeptical about behavioral evidence for it. The reason is that demonstrating decay behaviorally is exceptionally difficult. Proving that decay exists requires a demonstration that forgetting grows over time, in the absence of other activities such as the storage of new experiences or rehearsal. Rehearsal of the memory in question must

KEY TERMS

Interference: The phenomenon in which the retrieval of a memory can be disrupted by the presence of related traces in memory.

Trace decay: The gradual weakening of memories resulting from the mere passage of time.

be controlled because, as discussed earlier, retrieval strengthens memories, which would undercut efforts to see decay. As we will later discuss, storing new experiences after a trace has been encoded must be controlled because new memories introduce interference that might disrupt recall. When these constraints are considered, the person would essentially need to be kept in a mental vacuum, devoid of rehearsal, thoughts, or experiences that might contaminate the state of memory and complicate the interpretation of forgetting. To make matters worse, even if forgetting occurred in the absence of interference, it remains unclear whether the trace has become unavailable, or is merely inaccessible. Thus, it might be impossible to establish evidence for decay behaviorally, even if it does exist.

Correlates of time that cause forgetting

For the foregoing reasons, experimental psychologists have favored the view that time is merely correlated with some other factor that causes forgetting. Two possibilities have been examined. First, as time goes by, the incidental context within which we operate gradually shifts, perhaps impairing retrieval of older memories. Second, over time, people store many new similar experiences that might interfere with retrieving a particular trace. Although these factors do not disprove decay, they provide alternative explanations for the forgetting curve that do not rely on this process.

Contextual fluctuation

As discussed in Chapter 8, retrieval hinges on the number and quality of cues available during recall. When irrelevant cues are used, retrieval can fail. Retrieval can fail when a cue that was previously relevant changes over time. For instance, family members change in appearance, making them match less well the original cue associated to a memory. Moreover, when incidental context at retrieval does not match the one present at encoding, forgetting is more likely. One explanation of the forgetting curve then, is that as time progresses, changes in context become greater, on average because the world changes and we change. With time, we encounter new stimuli, people, and situations, and we have new thoughts and emotions. As such, one's incidental context will be most similar to the one that we were in a short while ago, and grow less similar over time. The idea that **contextual fluctuation** contributes to forgetting has been advocated in numerous models of memory (e.g. Mensink & Raaijmakers, 1988).

Changes in context may partially explain the striking phenomenon of *infantile amnesia*, which is discussed further in Chapter 12. Infantile amnesia refers to the difficulty most people have in remembering the first several years of their lives. One hypothesis is that the incidental contexts of young children are so vastly different from that of adults that childhood memories simply cannot be retrieved. Young infants have not developed language, for example, and so lack the conceptual and linguistic ambiance in which we exist as adults. Children also surely experience things differently from the adults that they become, with things seeming relatively larger at younger ages. So, infantile amnesia may rest in part on changes in environmental, cognitive, and perhaps emotional context.

Interference

Over time, experiences accumulate. Like the clutter of papers on your desk, adding new memories affects how easily we find things already stored. When memories are similar, this problem should be even worse, like having many similarly labeled papers on a desk. The idea that storing similar traces impedes retrieval is known as *interference*. Interference is likely

KEY TERM

Contextual fluctuation: The gradual drift in incidental context over time, such that distant memories deviate from the current context more so than newer memories, thereby diminishing the former's potency as a retrieval cue for older memories.

to be a serious issue when you consider that people are, by their nature, creatures of habit. People enjoy their routines, be they reading the newspaper in the morning, parking in the same spot each day, or getting their morning coffee at the same time. Sticking to routines, however, makes life less memorable. We remember what we had for dinner last night, but not 2 weeks ago. Such forgetting doesn't simply reflect the passage of time. We can easily remember experiences for a long time if they are unique: Having dinner at the neighbors' house a year ago is far more memorable than having dinner at our own house 3 months ago. It is the presence of other traces in memory that compromises retrieval. Because the number of similar traces will increase over time, interference provides a straightforward account of the forgetting curve. The emphasis on interference as a source of forgetting has a long history (Müller & Pilzecker, 1900) and was a preoccupation of research on memory for the nearly three-quarters of a century (see Postman, 1971; Crowder, 1976; Anderson & Neely, 1996, for reviews).

How does adding similar experiences into memory hurt us? To understand this, it is helpful to step back and discuss a fundamental discovery about what likely underlies interference. Early in the history of memory research, investigators identified a central feature in common to most situations associated with interference: Interference arises whenever the cue used to access a target (Figure 9.4, top left) becomes associated to additional memories. The

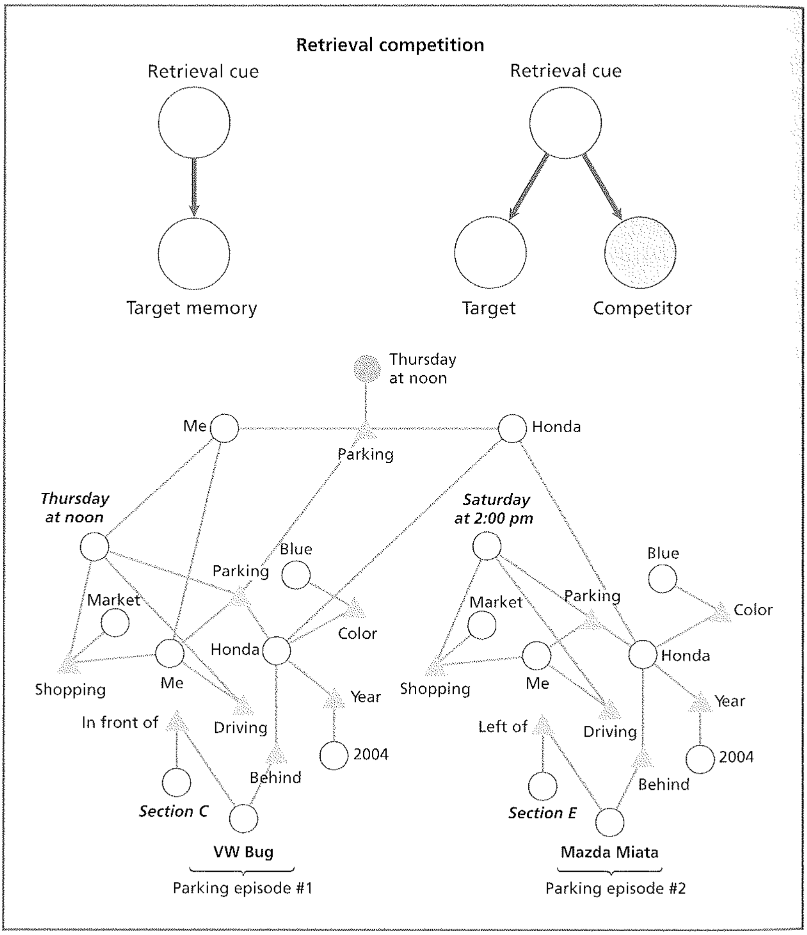

Figure 9.4 Top left: A retrieval cue associated to a single target item. Top right: A retrieval cue becomes associated to a competitor interfering with recall of the target. Bottom: A more complex example of interference, with multiple, shared retrieval cues and complex memories with many features. From Anderson and Neely (1996). Copyright © Academic Press. Reproduced with permission from Elsevier.

canonical interference situation is illustrated in the top right panel of Figure 9.4 in its most general form, with a single cue linked to many associates. By this view, progressing from a cue to a target depends not only on how strongly that cue is associated to the target but also on whether the cue is related to other items. Why does attaching more memories to a cue make retrieving a particular target difficult? Although theories vary about the particulars, most agree that when a cue is linked to multiple items, those items compete with the target for access to awareness, an idea known as the **competition assumption** (Anderson, Bjork, & Bjork, 1994). Essentially, a cue activates all of its associates to some degree, and they "fight" one another. As such, any associates other than the target memory are called *competitors*. In general, any negative effect on memory arising from having competitors is called interference. Interference increases with the number of competitors a target has. This idea is supported by the tendency for recall to decrease with the number of to-be-remembered items paired with the same cue, a generalization known as the **cue-overload principle** (e.g. Watkins, 1978). In essence, as a cue becomes attached to too many things, its capacity to access any one trace is compromised.

How do these ideas explain why storing similar memories causes interference? Consider an example in which you are recalling where you parked your car in a shopping mall you visit frequently. While parking, you will have encoded aspects of your parking experience into a memory. Other parking memories similar to this one will also contain characteristics of the target, including the fact that you drove a car, the type of car you drove (e.g. a 2004 blue Honda) and, perhaps, your goal of shopping. If important elements of the target (e.g. the concepts of yourself, of parking, and of your Honda) serve as the cues to your car's location, other memories sharing those features will be evoked as well. Figure 9.4, bottom, illustrates this by showing how the situation illustrated with one cue may be scaled up to the many cues available in this example (e.g. "Me," "Parking," and "Honda"). Thus, competition for a shared cue is a useful way of viewing interference between similar traces.

The notion of competition among items that share retrieval cues is very general. For instance, items in memory need not be episodes to compete. Indeed, even retrieving the meaning of a word can involve retrieval interference. To convince yourself of this, try the following demonstration, illustrated in Table 9.1. Each of the words listed in this table has entirely distinct verb and noun meanings, with the verb meaning being the less common. For each word, try to generate an associate of its verb meaning. For instance, for the word *duck*, you would generate a word like *crouch*, signifying that you thought of the verb meaning. Do this for each word as quickly as possible.

For native speakers of English, this task is perplexingly difficult because they instantly retrieve the noun meaning of the word, and must work to get past that dominant association. If this happened to you, you experienced competition from the noun meaning during the retrieval of the verb.

Interference phenomena

A number of qualitatively distinct situations produce interference. For instance, the storage of new experiences can interfere with retrieving older ones, but older memories can also impede retrieval of newer ones. In this section, we review some of the most important interference phenomena and key results that

KEY TERMS

Competition assumption: The theoretical proposition that the memories associated to a shared retrieval cue automatically impede one another's retrieval when the cue is presented.

Cue-overload principle: The observed tendency for recall success to decrease as the number of to-be-remembered items associated to a cue increases.

Table 9.1: Interference effects

	Cue	Related verb
E.g.	*Duck*	*Crouch*
1	*Loaf*	
2	*Post*	
3	*Court*	
4	*Root*	
5	*Sock*	
6	*Shed*	
7	*Fence*	
8	*Lobby*	
9	*Stump*	
10	*Fawn*	
11	*Lodge*	
12	*Sign*	
13	*Bark*	
14	*Pine*	
15	*Bowl*	
16	*Prune*	
17	*Duck*	
18	*Rail*	
19	*Sink*	
20	*Ring*	

Adapted from Johnson and Anderson (2004).

have been discovered. It is important to bear in mind that although the particulars of these situations vary, the underlying mechanisms that produce forgetting may in fact be similar. In the section to follow, we consider candidate mechanisms.

Retroactive interference

At the beginning of this chapter, you were asked to list all of the things you did yesterday, the same day last week, and the same day last year. If you did this exercise, you undoubtedly confronted the uncomfortable fact that you remember little of what has happened in your life. Why? As we have discussed, the difficulty might be due to several sources, including decay and contextual fluctuation. But there is an excellent chance that a lot of that forgetting comes about due to **retroactive interference.** Retroactive interference refers to forgetting caused by encoding new traces into memory in between the initial encoding of the target and when it is tested. Essentially, some process associated with storing newer experiences impairs the ability to recall ones farther back in time. With every new trip to McDonald's, every morning ride on the bus, and every day you spend seated in front of a computer screen at work, previous McDonald's trips, morning rides, and days at work grow farther from your mental grasp.

The methods used to study retroactive interference have tended to focus on simple materials that conform closely to the canonical interference situation described earlier. This phenomenon is often studied using the classic retroactive interference design illustrated in the left half of Figure 9.5. In the experimental condition, people study a first list of pairs (upper box), and then a second list. Very often the pairs in the first list (e.g. *dog–sky*) have their cue words repeated in the second list, but are paired with a new response word (e.g. *dog–rock*) that people have to learn in place of the older one. After the second list is learned, people are usually tested by giving them the first word of each pair and asking them to recall the response from the first list (e.g. *dog–?*). In the control condition, people also study a first list, but engage in irrelevant filler activity in the interval, during which people in the experimental condition study

KEY TERM

Retroactive interference: The tendency for more recently acquired information to impede retrieval of similar older memories.

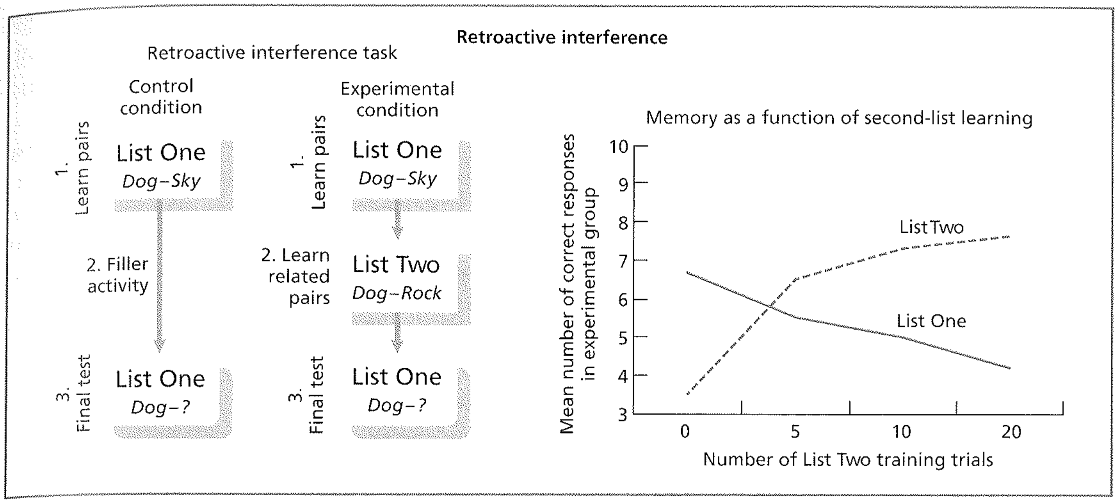

Figure 9.5 Left: A retroactive interference task in which participants learn two lists of word pairs, in series. A control group skips second-list learning. Right: Final cued-recall results (for both lists) by list-two training trials. Memory for list two increases with practice, while list one declines with list-two practice. Data from Barnes and Underwood (1959).

the second list. Thus, these two conditions allow us to ask the crucial question: "What is the effect of learning new information (i.e. list two) on the ability to remember information that was previously studied (i.e. list one), relative to a situation in which no additional information was learned at all (i.e. control condition)?"

The general findings are that: (1) introducing a highly related second list impairs the ability to recall items from the first list, compared to the control; and (2) increased training on second-list items continues to harm retention of first-list items further, as training progresses. This is especially true when the lists one and two share a common cue word (e.g. *dog*, as in the previous example); in fact, there is often little retroactive interference when the pairs on the two lists are unrelated. Thus, not every type of intervening experience impairs memory—the experience needs to be similar. A typical example of retroactive interference is illustrated in the right half of Figure 9.5, which is taken from a classic study by Barnes and Underwood (1959). Notice that as people were given increasing amounts of training on the second list of pairs, their memory for those pairs gets better, whereas their retention of first-list pairs grows quite a bit worse. We know that this increased forgetting is not due to the mere passage of time, because in the control condition the same amount of time has gone by in between learning the pairs and the final test. Thus, learning something new can impair memory substantially.

But are the lessons from artificial laboratory materials applicable to memory for personal experiences? It would be helpful if it could be shown that something like retroactive interference occurs with realistic memories. Such studies exist and generally confirm the importance of retroactive interference. In one study by Hitch and Baddeley, rugby players were asked to recall the names of the teams they had played earlier in the season (Baddeley & Hitch, 1977). Figure 9.6 shows the probability of their recalling the name of the last team played, the team before that, and so forth. It proved to be the case that most players had missed some games, either due to injury or other commitments, so that for one player the game before last might have taken place a week ago and for another it might have been 2 weeks or even a month before. It was therefore possible to ascertain whether forgetting depended on elapsed time or on the number of

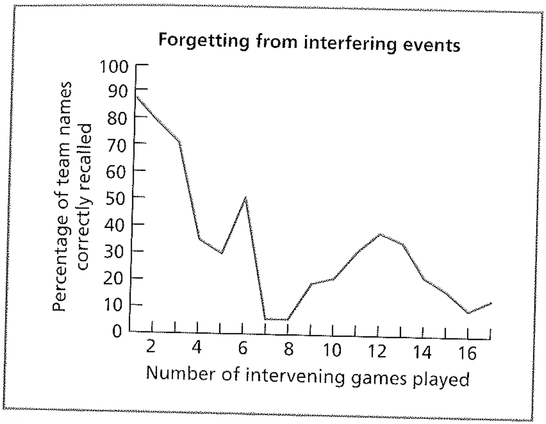

Figure 9.6 This graph, showing rugby players' memory for the names of teams recently played, demonstrates the tendency for recent events to interfere with memory of similar events from the past. Data from Baddeley and Hitch (1977).

intervening games. The result was clear. Time was relatively unimportant, whereas the number of intervening games was critical, indicating that forgetting was due to interference rather than trace decay. Apparently, their memory of having played a whole rugby game could be made less accessible simply because they have played many rugby games since then.

Proactive interference

One afternoon, I walked up to the top of an exceptionally steep street outside the psychology department to discover, to my horror, that my car had been stolen. After a moment's reflection, I realized that I had not parked my car there this *afternoon*, but rather this *morning*. This afternoon, I had parked my car on an entirely different, preposterously steep hill one street over. I was the unhappy victim of *proactive interference*, or the tendency for older memories to interfere with the retrieval of more recent experiences and knowledge. Most of us are acquainted with the irritation of proactive interference. It occurs, for example, when we fail to recall our new password because our old one intrudes during recall, refusing to be ignored or abandoned simply because it is out of date. Or, if we are seriously unlucky, we might call our current partner by our previous partner's name in an absent-minded moment. In each case, well-encoded events or facts rear their ugly head and disrupt retrieval of something more recent.

Although we have emphasized how retroactive interference affects long-term retention, proactive interference plays a powerful role in determining the rate of forgetting. This was demonstrated dramatically by Benton Underwood. Underwood (1957) was interested in explaining why participants who had learned a list of nonsense syllables should show so much forgetting after 24 hours. It occurred to Underwood that proactive interference was a real possibility. The reason was that almost all work on human learning at the time was done in a few laboratories, all of which used undergraduate participants. If you happened to be a student in one of these departments, you were likely to be required to participate for many hours in verbal learning studies. Underwood thought that it might be interference from the many *previous* lists of nonsense syllables that caused forgetting. Fortunately, it was possible to find out how many previous lists each participant had learned in other experiments and to plot the amount of forgetting in a 24-hour period as a function of this prior experience. In fact, naive students, who had no previous experience remembered 80% of the list items after 24 hours, whereas students with 20 or more prior learning trials on different lists remembered fewer than 20% 24 hours later. Proactive interference had a giant effect on retention, largely determining the rate at which students forgot the material after an extended delay.

Experiments examining proactive interference have often used an experimental design that is highly related to the retroactive interference design described earlier. The proactive interference paradigm (Figure 9.7) resembles the retroactive interference design, except that: (1) it tests people's memory for the list-two responses rather than the list-one responses; and (2) in the control condition, the rest period (or performance of irrelevant activity) replaces list-one learning rather than list-two learning.

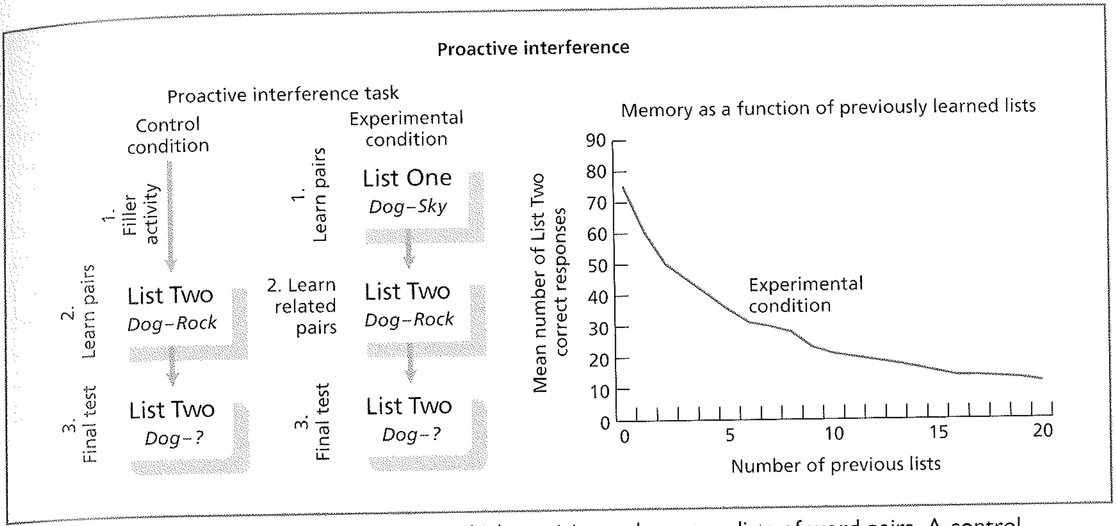

Figure 9.7 A proactive interference task, in which participants learn two lists of word pairs. A control group skips first-list learning. Right: A meta-analysis of final test cued-recall data following a 24-hour delay, given varying numbers of preceding lists. List-two memory decreases as the number of prior lists increases. Adapted from Underwood (1957).

Thus, this design allows us to explore how previously acquired knowledge (i.e. list one) might impair our ability to recollect new information (i.e. list two), relative to a situation in which the previous knowledge had not been learned (control, list two). Studies using the proactive interference procedure have demonstrated that people are more likely to forget items from a list when a prior list has been studied. The amount of proactive interference is greatest when the two lists share a common cue. Proactive interference effects are most severe when recall is tested rather than recognition.

Part-set cuing impairment

Recent exposure to one or more competitors exacerbates the problems we have in retrieving a target memory. For example, most of us have forgotten the name of someone and have been offered assistance by a well-meaning friend who supplies guesses about the name we are seeking. Unless the friend is lucky and guesses correctly, it often feels as though his or her suggestions make matters worse. Sometimes recall fails until a much later point when, unencumbered by the clutter of incorrect guesses, your mind yields the delinquent name. If you have had this happen, you have had firsthand experience with the phenomenon of **part-set cuing impairment**.

Part-set cuing impairment refers to the tendency for target recall to be impaired by the provision of retrieval cues drawn from the same set (e.g. category) of items in memory (Mueller & Brown, 1977). The basic finding was discovered by Slamecka (1968), who had people study lists composed of words from several semantic categories (e.g. *trees*, *birds*). On the final test, some people were given some of the members from each category as cues to help them recall the remainder; others were given no such cues. Of concern was people's recall of the remaining noncue items in the experimental condition relative to recall for those items when no cues were given.

KEY TERM

Part-set cuing impairment: When presenting part of a set of items (e.g. a category, a mental list of movies you want to rent) hinders your ability to recall the remaining items in the set.

Have you ever walked into a store, only to forget about your intended purchase? Blame part-set cueing, the tendency for the presence of some items as retrieval cues (like the CDs on display in the storefront) to impair one's ability to retrieve other items within the same set (the desired CD).

Slamecka expected that the cues would help recall for the noncue items. To his surprise, when recall was scored for the noncue items, people receiving cues performed worse than those who received no cues! This has become known as part-set cuing impairment because providing part of the set (in this case, part of the category) as cues impaired recall of the remaining items. Part-set cuing might be one reason why every musical CD that we make a "mental note" to purchase the next time we are at the store seems to disappear from our minds the moment we enter and peruse items displayed on the shelves.

The idea that supplying hints might impair memory is both surprising and ironic. In retrospect, however, it makes good sense, given the situation of interference described at the outset. Presumably, a set of items is defined by some common cue (e.g. *fruit* or *birds*), to which many items are associated. If presenting some items from the set strengthens their associations to the cue, perhaps stronger items provide greater competition during the retrieval of noncue items, impairing their recall. The idea that cues add more competition is consistent with the finding that as more members of the set are provided as cues, the worse memory becomes for the remainder (see Nickerson, 1984, for a review).

If people's instinct to be helpful and provide cues sometimes harms memory, what would happen if a group of people got together and tried to collaboratively remember things that they had all experienced or learned? Would one person's recounting prompt others to remember more than they would have, or might it cause part-set cuing impairment? Recent work indicates that when people get together to remember material that they each learned, they remember less when recalling the information as a group than they do when each person recalls information separately and their results combined into a common score. This phenomenon, known as **collaborative inhibition** can arise from the mechanisms that produce part-set cuing inhibition (Weldon & Bellinger, 1997). If group members are generating lots of items while you are listening, the interference this causes can disrupt your retrieval. Thus, research on part-set cuing might help to explain the effects of group effort on generating a diversity of new ideas and recollections.

Retrieval-induced forgetting

An ironic feature of human memory is that the very act of remembering causes forgetting. It's not that remembering harms memory for the retrieved experience. Rather, retrieval can harm recall of other memories or facts related to the retrieved item. Anderson et al. (1994) referred to this phenomenon as **retrieval-induced forgetting**.

KEY TERMS

Collaborative inhibition: A phenomenon in which a group of individuals remembers significantly less material collectively than does the combined performance of each group member individually when recalling alone.

Retrieval-induced forgetting (RIF): The tendency for the retrieval of some target items from long-term memory to impair the later ability to recall other items related to those targets.

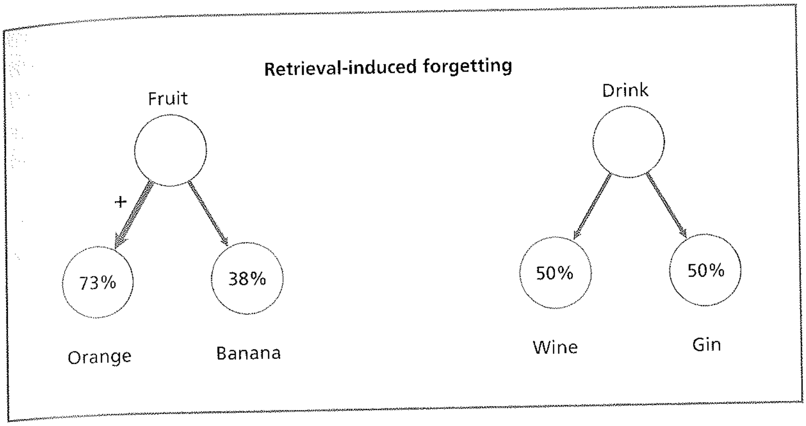

Figure 9.8 In this example, participants perform retrieval practice on *orange* but not *banana* or any members from the *drink* category (baseline). The final test scores indicate that, relative to baseline, practice facilitates recall of the practiced items, whereas unpracticed items from practiced categories suffer retrieval-induced forgetting. Adapted from Anderson (2003).

The phenomenon of retrieval-induced forgetting is usually studied with a procedure known as the **retrieval practice paradigm** (Anderson et al., 1994), which is illustrated in Figure 9.8. In this procedure, people first study simple verbal categories, like *fruits, drinks*, and *trees* for a later memory test. Following this, people are asked to repeatedly recall some of the examples that they just studied, from some of the categories. For example, participants might receive the cues *fruit-or—* for the retrieval of the item *orange*. Following retrieval practice, a test is given in which people are asked to recall all examples that they remember seeing from every category. Clearly, on this final test, people will recall the examples that they practiced quite well. More interesting, however, is how well they recall the remaining unpracticed examples (e.g. *fruit–banana*), compared to unpracticed items from baseline categories that are also studied, but none of whose examples receive retrieval practice (e.g. *drinks–scotch*). Strikingly, as can be seen in Figure 9.8, retrieval practice enhances recall of practiced items (e.g. *fruit–orange*), but it impairs related items (e.g. *fruit–banana*). So, it seems that the very act of remembering can cause forgetting.

If retrieval causes forgetting, students might have reason to be concerned about how they study for exams. Consider the plight of students who have limited time to prepare. You must prioritize your time, and the issue arises as to what to pass over. Research on retrieval-induced forgetting suggests that selectively reviewing facts impairs nonreviewed material, particularly related material. Neil Macrae and Malcolm MacLeod (1999) tested this idea by giving students facts like those they might learn in a classroom. Participants studied ten geography facts about each of two fictitious islands (*Tok* and *Bilu*; e.g. *The official language of Tok is French* or *Bilu's only major export is copper*). Students then performed retrieval practice. For one island, they practiced retrieving five of its ten facts. A final test followed, cued by the name of each island. Macrae and MacLeod found that practice facilitated the later recall of practiced facts (70%) over baseline facts about the unpracticed island (38%), but at the cost of impairing retention of related but unpracticed facts (23%). Similar effects have been demonstrated with complex textual materials when the test involved either short-answer or essay tests, although not on multiple-choice exams (Carroll, Campbell-Ratcliffe, Murnane, & Perfect, 2007). So, one must be careful about what one leaves out whilst studying, because omitting material will hasten its forgetting.

Selective retrieval occurs often in daily life. One situation of concern arises when members

KEY TERM

Retrieval practice paradigm: A procedure used to study retrieval-induced forgetting.

of law enforcement, detectives, and lawyers interview a witness after a crime. Answering interrogators, of course, requires retrieval. John Shaw, a psychologist who had once been a Los Angeles public defender, thought that such questioning might harm witnesses' memories for nonquestioned material, an intuition based on experiences with some of his own clients. To examine this possibility, Shaw, Bjork, and Handal (1995) told a group of participants to imagine that they had attended a party and that, on leaving, they noticed that their wallet was missing. Participants then watched slides of a student's apartment and paid attention to the details contained therein so that they might assist the police in an investigation. The slides contained a number of household items plus two categories of critical items (i.e. college sweatshirts and schoolbooks). Participants were then given structured questions about some of the objects (e.g. sweatshirts) during the *interrogation phase*. Consistent with Shaw's experience, they found that interrogating people about some stolen items impaired their memory for related items. So, retrieval-induced forgetting could have significant implications for how witnesses should be questioned.

If retrieval impairs memory, then simply discussing an experience with someone might alter whether people will remember what was omitted. Conroy and Salmon (2006) examined this idea by having young children participate in a staged event at school called *Visiting the pirate*, during which the children engaged in a number of activities across a variety of scenes. For example, in the *Becoming a pirate* scene, the children were asked to hoist a sale, bang a drum, put on pirate clothes, greet a pirate, and put their name in the pirate's book; whereas in the *Winning the key* scene, they might have fed a bird, looked through a telescope, steered the pirate ship, and done a dance. On the next 3 days, the children discussed the event with another experimenter, who asked them questions about only some parts, such as, "Tell me about the animal that you fed." On the final day, the children recalled the nondiscussed elements less well than did a control group of children who engaged in no discussion at all. Conroy and Salmon speculated that children's memory of their growing up years will be shaped by the way in which parents and family members reminisce, with nondiscussed aspects growing appreciably less accessible over time.

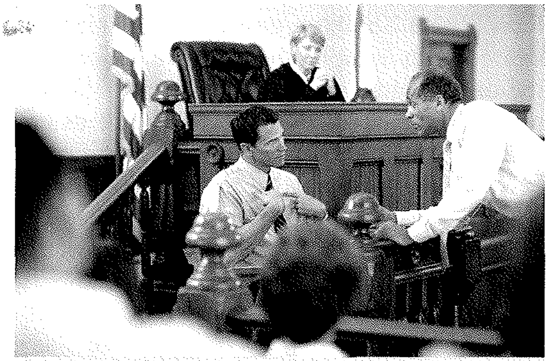

Retrieval-induced forgetting suggests that selectively reviewing facts impairs non-reviewed material. What implications could this have for the questioning procedures used, for example, in a court of law? © Tim Pannell/Corbis.

If discussions with other people about a shared past can lead one to forget what is not discussed, then forgetting can, in a sense, be contagious. A friend who has forgotten some parts of an experience will leave the forgotten parts out while reminiscing about it. Might selective remembering on the part of one person cause forgetting of the nondiscussed material in others? Alexandru Cuc, Jonathan Koppel, and William Hirst (2007) looked at this possibility in recent work on socially-shared retrieval-induced forgetting. One study replicated the experiment of Anderson et al. (1994; discussed earlier) with a twist: they had two people, seated side by side, studying the same pairs. In the retrieval practice phase, however, one participant performed retrieval practice whereas the other sat silently and observed, monitoring their partner's recollections for accuracy. Both then took the final test. As expected, the participant who performed retrieval practice showed retrieval-induced forgetting. Surprisingly, however, the silent observer also showed this effect. Cuc and colleagues observed the same effect when they used stories as materials; they even observed it when people were allowed to discuss the stories freely with one another: The nondiscussed elements of the story for one person were more likely to be forgotten by the other. It seems that when we are amongst others discussing past events, we spontaneously recall those events along with the person doing the recounting, and, in doing so, subject ourselves to retrieval-induced forgetting for whatever the speaker remains silent about. If this is so, then retrieval-induced forgetting might be one mechanism by which a society's collective memory of an event comes to be more uniform over time. It might also provide a means of political manipulation, when silences about certain facts or events are deliberate, and mass media is used to trumpet certain elements of the past. As Cuc and colleagues remark, "Silence is not always golden."

Retrieval thus appears to be a powerful force that shapes memory, for the better and for the worse. As discussed earlier, Marigold Linton's observations indicate that retrieval greatly enhances the longevity of a memory, but the current results show that, when retrieval is incomplete, the benefits might be offset by the forgetting of other things. To understand the importance of this finding, one need only consider how pervasive this basic process is in our daily cognitive experience. That is, any cognitive act that makes reference to traces stored in memory, which is likely to be all processes, employs retrieval. If retrieval is a source of forgetting, then accessing what we already know might contribute to forgetting, independent of the encoding of new experience. The role of retrieval in causing forgetting has led to a new perspective on why the situation of interference is associated with forgetting (Figure 9.9). We discuss this perspective shortly.

Interference mechanisms

As the preceding discussion illustrates, many "interference" situations impair retention. Although these phenomena describe *when* forgetting will arise, they do not say *how* forgetting occurs—that is, they do not specify the mechanisms. Why does presenting cues impair recall? Why does retrieval-induced forgetting occur? Why does introducing new learning impair retention of previously acquired material? First, we consider classical mechanisms proposed to explain interference, and show how they can be extended to explain phenomena like part-set cuing and retrieval-induced forgetting. Then we consider a more recent view in which inhibitory processes associated with retrieval cause forgetting.

Associative blocking

Once, while recalling the British term for what Americans call a "Christmas ornament," I persistently recalled "Christmas balls" (what

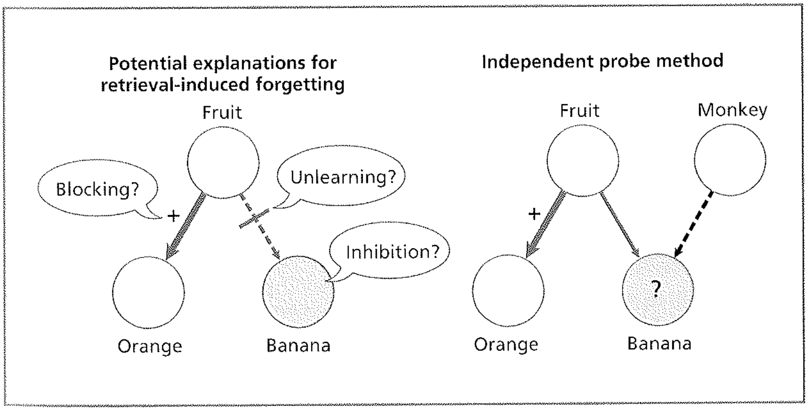

Figure 9.9 Left: Practiced items could block target recall during the final test, producing retrieval-induced forgetting. Alternatively, the connection between cue and target could have been unlearned during retrieval practice. Inhibition involves a reduction in the target memory's activation, itself. Right: Only inhibition correctly predicts that forgetting generalizes to independent cues.

Belgians call Christmas ornaments), instead of "Christmas baubles." "Christmas balls" kept intruding until I gave up. In essence, "Christmas balls" blocked "Christmas baubles." After drifting off to other activities, the right answer "popped" to mind. Perhaps something like this tip-of-the-tongue experience might explain interference. For instance, in retroactive interference, people might forget first-list responses because the cue used to access them now elicits the second-list responses. In part-set cuing, presenting exemplar cues might strengthen their association to the category, leading them to intrude when people try to retrieve noncue exemplars. In each case, a cue elicits a stronger competitor, leading us to perseverate helplessly on something that we know to be incorrect. The idea that such a process explains interference was proposed by McGeoch in his (1942) *response competition theory*, modern versions of which are known as **associative blocking** (see Anderson et al., 1994).

The core assumption of the blocking hypothesis is the idea that memories compete for access to awareness when their shared cue is provided. The degree of interference should increase as the cue grows more strongly associated to the competitor, exhibiting what Anderson et al. (1994) refer to as *strength-dependent competition*. But how does a stronger competitor impair recall? Consider an example from retrieval-induced forgetting, in which you are trying to recall *banana* after having practiced *fruit–orange*. According to the blocking theory, the cues on the final test (e.g. *fruit*) to recall *banana* lead the person to accidentally retrieve the stronger practiced item, *orange*. Once accidentally retrieved, *orange* will achieve greater prominence, having been practiced again, making it even more likely to be accidentally retrieved. And so the cycle would continue, because, with each accidental retrieval, the wrong answer grows stronger. Eventually, people might simply give up. So, according to the blocking theory, people forget unpracticed exemplars of practiced categories because associations to the practiced memories dominate retrieval. Blocking can also explain the cue-overload principle: The more memories associated to a cue, the more likely it should be to accidentally retrieve a wrong answer, setting the blocking process in motion. If so, the reason you can't remember your dinner 4 months ago is because retrieval calls to mind recent dinners to such an extent that you give up.

Associative unlearning

Suppose that your acquaintance describes a conversation at a party several years ago. You might recall, in good detail, elements of the party, including your friend's attendance, various conversations, as well as several amusing events. However, you might forget discussing a topic with your friend, despite your friend's most confident confirmations—even when you clearly recollect discussing the topic. Subjectively, it seems as though your memory has become fragmented, impairing your judgment about how elements of the experience go together. This apparent fragmentation can reflect damage to the associations between elements of that event caused by storing subsequent experiences. Perhaps such damage underlies my inability to recollect knocking over the Christmas tree.

Research on the **unlearning** hypothesis of retroactive interference (Melton & Irwin,

KEY TERMS

Associative blocking: A theoretical process hypothesized to explain interference effects during retrieval, according to which a cue fails to elicit a target trace because it repeatedly elicits a stronger competitor, leading people to abandon efforts to retrieve the target.

Unlearning: The proposition that the associative bond linking a stimulus to a memory trace will be weakened when the trace is retrieved in error when a different trace is sought.

1940) is relevant to these ideas. According to the hypothesis, the association between a stimulus and a trace will be weakened whenever that trace is retrieved inappropriately. In effect, the bond between the cue and the target gets "punished." For example, suppose that you try to retrieve the new password to your e-mail account. According to the unlearning view, if you recall your old password and realize the mistake, the association between the cue password and the original password will get weakened, decreasing the chances that it will pop up again in the future. If the old password is punished often enough, the association might grow so weak that it will no longer activate that trace; the stimulus will be decoupled from the response. This view can explain retrieval-induced forgetting if we assume that, during retrieval practice, competing items intrude and are punished. It explains retroactive interference in the same way. So, whereas blocking attributes forgetting to very strong competitors, unlearning says that associations into the target are too weak.

The unlearning and blocking hypotheses are not incompatible. In fact, according to the classical two-factor model of retroactive interference (Melton & Irwin, 1940), both mechanisms are needed. It is worth emphasizing, however, that proof of unlearning is difficult to establish, for the same reasons that it is difficult to prove that memories are permanently forgotten, as discussed earlier in this chapter. Also, although blocking explains why forgetting appears to grow in magnitude as competitors are strengthened, there are reasons to doubt whether strengthening a competitor, by itself, produces forgetting, as will be illustrated shortly. For these reasons, an alternative view has arisen that attributes the forgetting arising from interference to inhibitory processes.

Inhibition as a cause of forgetting

The preceding discussion raises an important point: Sometimes it is maladaptive for a trace to be accessible. Retrieving a target can be disrupted by highly accessible competitors and people need a way to limit this distraction. Although unlearning is one way to accomplish this, another is to inhibit the offending trace. Consider an analogy. Suppose that you normally wear a watch, but one day the wristband breaks and you can't wear it. If someone then asks you what time it is, you might look at your wrist reflexively, even when you know the watch is absent. This might happen several times before you learn to look at the clock instead. Clearly, what is normally a useful and overlearned habit has, for the time being, become an inappropriate response that must be shut down so that an alternative but more appropriate response can be given. Humans and other organisms have the ability to terminate responses, either so that an alternative can be given or so that all responding can be stopped. Stopping is thought to be accomplished by a mechanism that inhibits the response. Inhibition reduces the activity level of the response, ceasing its production in a manner analogous to how inhibiting a neuron would reduce its influence on other neurons.

The same demands confronted in shutting down interfering responses occur for internal actions, such as retrieval. As discussed in the section on proactive interference, if somebody asks for our telephone number, we might automatically remember our old number even though we have switched phones. Recalling the new number requires that we stop retrieval of the old one, which can be accomplished by inhibition. If the old number is inhibited, however, it will grow harder to recall, even if it remains available. In the context of retrieval-induced forgetting, *banana* might become activated and intrude during the retrieval of *fruit–or—*. To facilitate the retrieval of *orange*, perhaps *banana* is inhibited, with persisting inhibition making it harder to retrieve that item. *Banana*, like the habit of looking at one's wrist, might be inhibited to support the current goals. Are inhibitory processes engaged during retrieval?

One area of research that has examined the role of inhibition during retrieval is retrieval-induced forgetting (Levy & Anderson, 2002). Inhibition makes several predictions about retrieval-induced forgetting that are not made by either the blocking or unlearning theories. According to inhibition, performing retrieval practice on *fruit–orange* impairs the recall of *banana* because *banana*, as the competing memory (like looking at the wristwatch), is inhibited by activation-reducing mechanisms. If *banana* is truly inhibited, one might imagine that *banana* would be harder to recall generally, whether one tests it with *fruit* as a cue, or, say, another unrelated associate, such as *monkey–b—*. In other words, inhibition predicts that retrieval-induced forgetting should generalize to new cues, thus exhibiting *cue independence*. By contrast, both blocking and unlearning attribute forgetting to problems with the associations linking *fruit* to either *banana* or *orange*. Hence, according to these theories, retrieval-induced forgetting should be cue dependent. That is, as long as you switch to another cue, like *monkey*, which circumvents the stronger association from *fruit* to *orange* and the potentially weaker association from *fruit* to *banana*, no impairment should be found for *banana*. Cue-independent forgetting has been observed many times (Anderson & Spellman, 1995; see Anderson, 2003, for a review), indicating that inhibition does play a role in causing retrieval-induced forgetting.

According to the inhibition hypothesis, the need to overcome interference during retrieval triggers inhibition. If so, then active retrieval on practiced items should be necessary to induce forgetting of competitors. For example, simply replacing retrieval practice trials (e.g. *fruit–or—*) with a chance to re-study *fruit–orange* multiple times should eliminate later forgetting of competitors like *banana*. Forgetting should disappear because giving people *fruit–orange* to study eliminates any struggle to retrieve *orange*, and thus any need to resolve interference from *banana*. This retrieval-specificity property is a consistent feature of retrieval-induced forgetting (see Anderson, 2003). Thus, even though both retrieval practice and extra study exposures strengthen memory for the practiced items to the same degree, only retrieval practice impairs retention of the unpracticed competitors. There appears to be something special about the need to reach into memory and retrieve something that induces forgetting, consistent with the idea that inhibition is involved. This finding doesn't favor the blocking hypothesis, however, which predicts that strengthening practiced shapes should impair recall of competitors, regardless of whether strengthening is accomplished by retrieval or study.

If inhibitory processes overcome interference from competitors, the amount of retrieval-induced forgetting should depend on the degree of interference during retrieval practice. If the other associates of a cue don't cause interference, inhibition should be unnecessary. In an early demonstration of this, Anderson et al. (1994) varied whether competing items were high-frequency examples of their respective categories (e.g. *fruit–banana*) or low-frequency examples (e.g. *fruit–guava*). Intuitively, one might imagine that a high-frequency example like *banana* would be resistant to forgetting, whereas a low-frequency item might be vulnerable. The analogy to the wristwatch example, however, suggests that the opposite could be true. It is precisely because one reflexively checks one's wrist when the watch is not there that one must inhibit that response to prevent it from occurring. If so, then high-frequency examples, like *fruit–banana* might be prime targets for inhibition because they come to mind so readily, whereas low-frequency exemplars might not need to be inhibited. This is exactly what Anderson and colleagues found. This property is known as *interference dependence*, or the tendency for retrieval-induced forgetting to be triggered by interference from a competing item.

One important feature of retrieval-induced forgetting that speaks against blocking theories is that the amount of forgetting appears unrelated to how strong the practiced associations become as a result of practice. Research on retrieval specificity, for example, shows that it is possible to greatly strengthen practiced items through repeated study, without impairing unpracticed competitors. If strengthening practiced items were enough to cause forgetting, one should have observed forgetting under these circumstances. Indeed, strengthening a competitor may not be necessary at all to trigger retrieval-induced forgetting. In a recent study, Benjamin Storm, Elizabeth Bjork, Robert Bjork, and John Nestojko (2006) had the clever idea to see whether it was merely the retrieval attempt that created retrieval-induced forgetting. Participants in this retrieval practice paradigm were, for some categories, given retrieval practice cues that were impossible to complete. So, for example, they might have received the cue *fruit–lu—* to complete during retrieval practice, even though no fruit begins with *lu*. Strikingly, even though people could not complete any of these retrieval practice tests, they showed as much retrieval-induced forgetting for the remaining unpracticed exemplars as they did for categories in which retrieval practice trials could be completed. So the struggle to extract a trace from memory, in the face of interference is the important trigger for retrieval-induced forgetting, not the strengthening of the practiced items. This property is referred to as *strength independence*.

Taken together, the properties of cue independence, retrieval specificity, interference dependence, and strength independence converge to support a role of inhibition as a source of forgetting (Table 9.2). If so, it suggests that many of our experiences with forgetting might arise from the need to control interference. It is precisely because we are distracted by momentarily irrelevant information in our memories—those unintended looks at our "mental wristwatch"—that we engage inhibition to refocus on what we hope to retrieve from memory. On the one hand, it might seem ironic that the mechanisms

Table 9.2: Properties of retrieval-induced forgetting

Property of retrieval-induced forgetting	Description
Cue independence	The tendency for forgetting caused by inhibition to generalize to novel test cues on the independent probe test (e.g. *monkey–b—* for *banana*, which was originally studied with the cue *fruit*)
Retrieval specificity	Active retrieval from long-term memory is necessary to induce forgetting of related information. For example, having to retrieve *orange*, given *fruit–or—* generates retrieval-induced forgetting of unpracticed competitors (e.g. *banana*), whereas simply studying the intact pairing (*fruit–orange*) does not
Strength independence	The degree to which competitors are strengthened by retrieval practice is unrelated to the size of the retrieval-induced forgetting deficit. Thus, strengthening an item by presenting the intact pairing (*fruit–orange*) does not induce retrieval-induced forgetting, whereas engaging in an impossible retrieval attempt (e.g. *fruit–lu—*) still results in forgetting of unpracticed competitors
Interference dependence	Interference by competitors during retrieval of targets is necessary for retrieval-induced forgetting of those competitors to occur. Therefore, high-frequency competitors (e.g. *fruit–banana*), which pose greater competition than low-frequency competitors (e.g. *fruit–guava*) are more likely to be inhibited than vice versa

we use to direct retrieval are the ones that ultimately contribute to forgetting. On the other hand, as Robert Bjork suggests, such forgetting might be adaptive because it helps to reduce interference from information that might no longer be as relevant as it once was (Bjork, 1989). If information remains in memory and can be revived (e.g. by re-exposure), forgetting may be very functional.

A FUNCTIONAL VIEW OF INCIDENTAL FORGETTING

Experimental psychologists have traditionally focused on passive mechanisms of forgetting, including trace decay, contextual fluctuation, the use of inappropriate retrieval cues, and interference processes such as blocking. The presumption has been that people are the passive victims of forgetting, with memory loss arising from factors that simply happen to us, such as random changes in the environment, and the addition of traces into memory. Although such processes contribute to forgetting, research on inhibition suggests a different outlook. According to the inhibitory control view, much of the forgetting that we experience arises from the need to control the retrieval process in the face of competition. It is the process by which we combat interference—inhibition of competing traces—that precipitates forgetting, not the mere presence of other traces in memory. By this view, reducing the accessibility of competing traces is adaptive because it facilitates retrieval, but also because it makes subsequent retrievals of the same information easier, reducing future competition. These observations highlight that inhibitory processes can be quite adaptive. This functional view conceptualizes forgetting as a positive outcome, and highlights how a properly functioning memory system must be as good at forgetting as it is at remembering (Bjork, 1988; Anderson & Spellman, 1995; Anderson, 2003; Bjork, Bjork, & Macleod, 2006). Thus, rather than being victims of forces beyond our control, many cases of forgetting might be tied to the very mechanisms that enable the effective control of cognition.

SUMMARY

Whether we like it or not, the vast majority of life's experiences are forgotten. What mechanisms underlie forgetting? To address these questions, we must first understand the characteristics of forgetting, including the scope of the problem and what constitutes forgetting. The fundamental fact is that forgetting increases as time goes by. Research on the forgetting curve reveals that the rate of forgetting is precipitous at first, but then tapers off gradually, following a logarithmic form. Many types of material follow this pattern, although the rate of forgetting varies depending on the nature of the test, and how well learned material is. Recall is more difficult than recognition, with the latter revealing evidence of retention not observed with recall. Forgetting is dramatically reduced for well-learned facts and experiences, especially when materials are repeatedly retrieved.

Experimental psychologists distinguish between the accessibility of an item and its availability, with loss of availability referring to the elimination of an item from storage. Both deficits in accessibility and availability are considered forgetting. It is difficult ever to

establish that a trace has been permanently lost, however, as it is always possible to suppose that the right retrieval cues have not been found. Thus, although permanent forgetting might exist, and might even be likely when one considers the biological system on which memory rests, it remains difficult to establish a case for this possibility through behavioral evidence alone.

To understand the fundamental fact that forgetting increases over time, several approaches have been offered. The most straightforward approach proposes that memories get weaker over time. Although decay is likely to be a biological reality and some evidence exists in non-humans, establishing a case for decay at the behavioral level is difficult because one has to control for the influence of other known forgetting factors, and establish that memories have been permanently lost. A different approach is to suppose that it is not time, *per se*, that causes forgetting, but some factor correlated with time, such as contextual fluctuation or interference. Both of these theories can accommodate the basic features of the forgetting curve, while also explaining mysteriously long-lived memories that would be difficult to explain if decay were the primary cause of forgetting.

The idea that similar memories interfere with one another during retrieval has occupied a central position in theories of forgetting. The fundamental situation of interference arises when a retrieval cue becomes associated with multiple traces. When a cue is presented for the retrieval of a target, the other associates of that cue compete with the target for access to awareness. The more competitors that are attached to a cue, the worse recall of any one item becomes, a generalization known as the cue overload principle.

Many interference phenomena have been discovered. For example, previously encoded traces can disrupt retrieval of recently learned ones, a finding known as proactive interference. Conversely, recently encoded traces can impede the retrieval of older memories, known as retroactive interference. Both proactive and retroactive interference are believed to contribute to the increase in forgetting over time, insofar as time is correlated with the storage of similar traces. In part-set cuing impairment, presenting some material as cues for the retrieval of other related memories impairs retention, showing that even after traces have been stored, recent exposure to some impairs recall of others. Research also indicates that retrieval impairs retention of related traces, known as retrieval-induced forgetting.

The mechanisms that underlie interference have long been a subject of research. Blocking theories attribute interference to the tendency for stronger traces persistently to intrude during retrieval of weaker ones, leading the person to abandon search. Unlearning theories propose that interference causes destructive changes to the associations that underlie a trace, as a result of learning mechanisms that punish inappropriate retrievals. Inhibition theories propose that forgetting arises, in part, from the suppression of interfering traces by inhibitory mechanisms that resolve competition. Research using the retrieval induced forgetting paradigm has provided specific evidence favoring the inhibitory view. This research suggests that some of the forgetting we experience arises as an adaptive consequence of controlling the retrieval process. This view highlights, more broadly, the notion that forgetting may sometimes be desirable. As the case of AJ illustrates, having a perfect memory can be a burden. We take up this topic further in Chapter 10 on motivated forgetting.

FURTHER READING

- **Anderson, M. C., & Neely, J. H.** (1996). Interference and inhibition in memory retrieval. In E. L. Bjork & R. A. Bjork (Eds.), *Memory: Handbook of perception and cognition* (2nd edn). (pp. 237–313). San Diego: Academic Press.
- **Levy, B. J., & Anderson, M. C.** (2002). Inhibitory processes and the control of memory retrieval. *Trends in Cognitive Sciences*, 6, 299–305.
- **Smith, S. M., & Vela, E.** (2001). Environmental context-dependent memory: A review and meta-analysis. *Psychonomic Bulletin and Review*, *8(2)*, 203–220.